A Comprehensive and Cost-Effective Computer Infrastructure for K-12 Schools

Gary P. Warren and Jeffrey M. Seaton

A White Paper from NASA Langley Research Center's HPCC/IITA K-12 Program

1.0 Introduction

The integration of computers and telecommunications into primary and secondary education has become a high priority for many school districts. A wealth of projects are currently being funded by government entities and private industry to bring technologies, such as the Internet, into the K-12 schools. Building a computer infrastructure capable of supporting popular Internet activities, such as obtaining information via the World Wide Web (WWW), is not a large technical challenge if monetary and support resources are available. Several school systems have been awarded large grants to demonstrate the effectiveness of using both the Internet and computational science to improve curricula. The challenge, however, is to build and support a computer infrastructure that satisfies the needs of teachers and administrators, is compliant with the Internet Engineering Task Force (IETF) recommendations[1], and is affordable within the school system's current technology budget. Unfortunately, most school districts that wish to build a computer infrastructure in accordance with the IETF recommendations feel they must seek outside funding sources to cover the high capital and recurring cost items.

The High Performance Computing and Communications (HPCC) Program is a federal program that has many diverse projects primarily focused on keeping the United States at the forefront of computer technology. One of these programs, the Information Infrastructure Technology and Applications (IITA) K-12 Program, is an educational outreach activity that is designed to enhance science and math curricula in the K-12 community. However, the use of computational science and telecommunications to enhance curricula requires a substantial computer infrastructure. As part of the IITA K-12 Program, NASA Langley Research Center has been developing a computer infrastructure model with low recurring costs, to meet the demands of large school networks. The initial testbed for this low-cost computer infrastructure model included six public high schools from six different school districts and one consortium school that represented the same six school districts. This pilot program, called the HorizonNet project and shown in figure 1, has been expanded by the participating school districts to include all school buildings on the Virginia Peninsula.

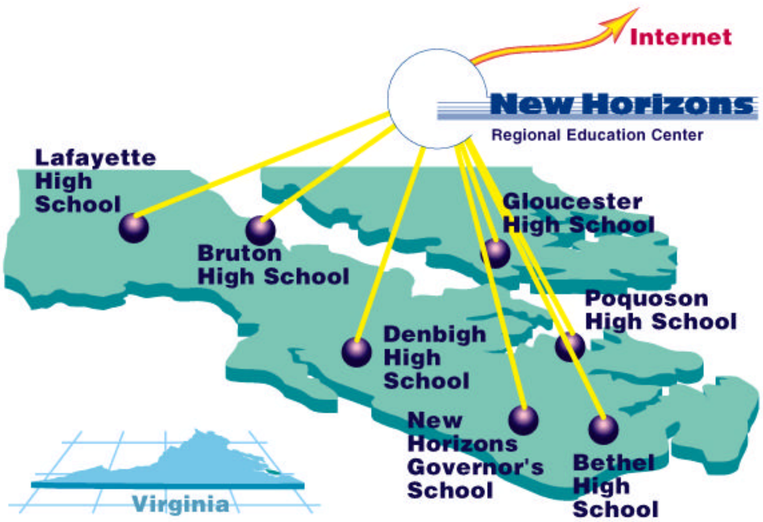

FIGURE 1. Pilot Schools Involved in the HorizonNet Project

New Horizons Regional Education Center is supported by six school districts. This consortium allows the six school districts to consolidate high-cost specialized programs for the region. Building a computer infrastructure with full Internet connectivity is costly, and individual school costs can only be minimized by using a regional approach.

The key to developing a successful computer infrastructure is to ensure that the foundational architecture will support Internet access for an entire school. The infrastructure developed here does not support a single or a few machines that dial into a bulletin board system (BBS) or access the Internet by using terminal emulation but allows up to 253 machines per school building to independently access all functionality associated with the Internet. As the demand for faster Internet access grows in the school, the model has a straightforward expansion path that does not require replacing the system and retraining personnel. The functionality provided by this model is identical to that which exists at most universities and government laboratories. By combining the resources of school districts in other regions, the computer infrastructure model represented in the HorizonNet project can be implemented on a much larger scale.

This paper addresses many issues associated with building a computer infrastructure with minimal cost while creating a solid foundation that can support curricula development. Although no single solution exists for providing Internet connectivity, the HorizonNet project demonstrates what is possible with current low-cost technologies. The HorizonNet project has demonstrated that an entire network of computers can be driven with acceptable performance over a standard telephone line. This paper details an implementation plan that has been successfully used by many schools in the southeastern region of Virginia. Because no one implementation plan will work with every school, this plan may require modifications to meet individual school district needs. However, all implementation plans should follow the guidelines for K-12 networking as outlined by the IETF. The computer and network model presented here has been shaped by the needs of teachers as well as by the reality of budget constraints; most importantly, it follows the guidelines established by the IETF.

2.0 Overview

2.1 Functionality

The functionality of any computer located in a school and connected to the network is far greater than is generally perceived. Any computer that is connected to the network and is configured to directly send and receive Internet traffic (TCP/IP) can run Internet applications locally and access information globally, as depicted in figure 2.

FIGURE 2. Access to the World Wide Web

This configuration has several advantages over a dial-in type of access, such as text-based terminal emulation. First, the requirement for an individual phone line for each computer that is accessing the Internet is eliminated. Second, Internet applications have the same look and feel of other applications that run locally on the computer. Userfriendly applications are available for computers that are running operating systems based on the Apple Macintosh and Microsoft Windows. The key issue is not the operating system that is used on individual desktop computers, but system-wide network language compatibility. If the computer can speak TCP/IP, then it can communicate with any other computer on the local area network (LAN) that also speaks TCP/IP, as well as with any other computer on the Internet.

2.2 The Building Blocks

Connecting every school to the Internet is an obvious national priority. However, serious considerations are associated with this connectivity. The Internet was originally conceived in government laboratories and universities as a research tool. Note, however, that a tremendous computational infrastructure already exists in major universities and government laboratories in comparison with what one would find in a typical public school district. Even commercial businesses that have connected to the Internet and have successfully taken advantage of its resources have realized the tremendous commitment, in terms of both staffing and financial investment, that is necessary to maintain this resource. Add to this commitment the lack of current curricula that are designed to incorporate a state-of-the-art computer infrastructure, and justification of the expense in the K-12 environment becomes difficult.

Because of the pressure placed on schools to provide Internet connectivity, typically paired with no additional funding to do so, schools sometimes fall victim to vendors who promise high-performance systems without high recurring costs. These systems are typically proprietary in nature and are difficult to integrate well with the current open standards that have evolved in large universities and government laboratories. Additionally, multiple vendors are typically needed to adequately address all areas of expertise shown in figure 3.

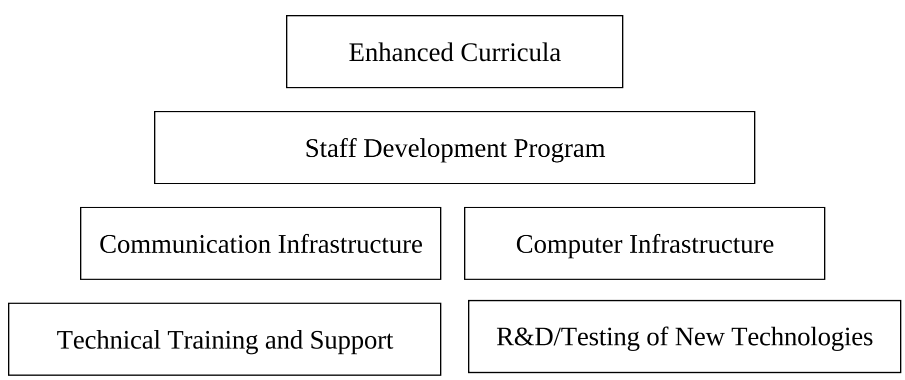

FIGURE 3. Required Building Blocks for Enhancing Curricula with Computational Science

In order for a computer and communication infrastructure to be successful, each building block represented in figure 3 must be carefully addressed. If the implementation model presented here is modified, each building block must be addressed in similar detail. The success of the HorizonNet project is based on the following key elements:

- Regional approach
- Standardized communication infrastructure
- Maximized performance using innovative technologies
- Leveraging the Internet server
- Distributed technical training and support

3.0 Regional Approach

High-cost programs can be made more affordable to communities through regional cooperation. Internet connectivity, including all required elements, is an extremely high-cost venture. The model presented here brings multiple school districts together within a consortium. This consortium is responsible for providing the high-cost items such as required personnel and equipment, establishing a high-quality network operations center (NOC), providing high-speed access to the Internet, and configuring routers and Internet servers, for which most school districts typically do not maintain the in-house expertise.

For any type of Internet connectivity, the recurring costs are usually of greater significance than the initial capital investment in equipment. Minimizing these recurring costs is one of the main goals of the model presented here. Listed in Table 1 are the basic elements associated with full-time connectivity, along with the approximate annual cost of these items in the HorizonNet project. This table includes all items, including communication links, support personnel, and maintenance of the Internet connection, which are associated with the potentially large networks that school districts represent.

TABLE 1. Responsibilities and Recurring Costs of Regional Center

Regional Center Responsibility	Approximate Annual Cost
Provide T1 leased line to Internet provider	$12,000
Establish routing for all school building networks (full class C network per building) with Internet provider	$24,000
Full-time UNIX/network administrator	$65,000
Equipment and software maintenance	$10,000

Each item listed in Table 1 is covered in greater detail later in this paper; this summary is designed to summarize the high costs associated with Internet connectivity when all elements are addressed.

One of the major costs shown in Table 1 is the personnel cost. Although often overlooked in any given model, the complexities associated with both network support and the required infrastructure demand that qualified personnel be assigned to maintaining the infrastructure. The model presented here minimizes personnel cost by standardizing the computer infrastructure located within each building. Standardization does not mean that every computer within a building or district must be manufactured by only one vendor, such as Apple or Intel, but that the supporting computer infrastructure that drives the network must be standardized.

This model uses distributed network servers based on the UNIX operating system, which has been in use for over 20 years and offers a stable environment for providing network services. Most of the software packages used to provide Internet services for large networks currently use computer platforms based on the UNIX operating system. With the explosive growth of the Internet, other operating systems are now offering the services required for TCP/IP networks. However, the UNIX operating system has the advantage of years of testing. Additionally, the preemptive multitasking and multi-threaded capabilities of the UNIX operating system allow many applications to be simultaneously executed while maintaining high performance and stability. The down-side to using the UNIX operating system is the knowledge required to set up and maintain the system. However, the cost of a UNIX system administrator can be reduced by using the regional approach for support.

Each school building that connects their network to the consortium must purchase a UNIX-based system that is acceptable to the consortium. Note that a particular brand of computer system should be standardized for the region to minimize the additional support time required. For example, the testbed associated with the New Horizons Regional Education Center was developed with Sun Microsystems Computer Corporation's SPARCstations. Although the model outlined here can be built with other vendors, a mixture of various UNIX implementations within the consortium will increase support costs.

Figure 4 shows the annual recurring costs for each connected school building that utilizes the HorizonNet model. As shown, when only a small number of schools are supported, the costs are fairly high; however, the costs become reasonable as the number of participating schools increases. The spike in the graph shown in figure 4 results from hiring additional support personnel after 40 school buildings are connected.

4.0 Standardized Communication Infrastructure

The creation of a communication infrastructure based upon open-standards is important in today's Internet-based environment. A properly constructed communication infrastructure allows an individual to access information outside the computer he or she is using. For school networks, scalability is very important. School networks, if used by students, can potentially be large networks; thus, from the outset, any communication infrastructure must be expandable to cover additional students and administrators without replacing equipment and retraining personnel.

Any given network includes two primary segments:

1. Communication lines
2. Communication protocols

Networking is more than simply obtaining a fast communication line from the local phone or cable company. Although the tendency is to concentrate on establishing fast communication lines, the establishment of the language the computers use to talk to one

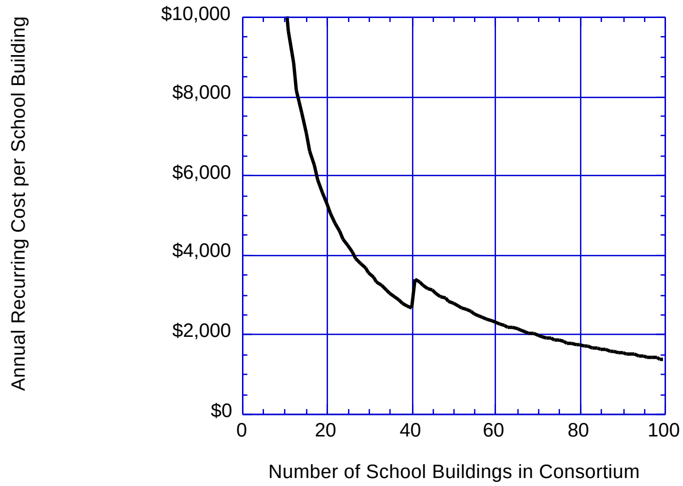

FIGURE 4. Sharing the Recurring Costs of Internet Connectivity

another across these communication lines is more important. Efficient use of the communication lines is also important. For example, downloading the same file from the Internet hundreds of times is not an efficient use of the communication line. The efficiency of various protocols relative to the size of a network can be used to determine the bandwidth required. The TCP/IP suite of protocols was used in this program after these factors, including the IETF recommendations, were taken into account.

Many popular networking methods that perform well for small school networks do not scale well to support thousands of simultaneous users. This section addresses some issues associated with creating a computer infrastructure that is comprehensive, scales well, and has a relatively low start-up cost. The goal of this model is to maximize performance at the desktop level. That is, performance of a communication infrastructure should be measured at an individual's desktop computer (i.e., the speed at which a user can send and receive data) and not on the sole fact that a high-speed (or low-speed) connection to the outside world is in place. If a school building has a fiber-optic connection to the Internet, but each user must wait 3 minutes for each e-mail request to be processed, a serious design problem exists in the network model.

4.1 Local Area and Wide Area Networks

A distinction needs to be made between the idea of a LAN and that of a wide area network (WAN). Although both are used to connect computers, the technologies employed for LAN's are much different than technologies employed for WAN's. A LAN is the local network within a given school building, whereas a WAN is the network that connects the school building LAN's together. High-speed access between computers is generally employed with LAN's because the distances are relatively short. Conversely, lower speed lines are commonly used with WAN's because the distances are substantially longer.

This paper does not address the creation of a well-designed LAN, other than to give the basic recommendations of the IETF:

1. An Ethernet network (e.g., 10base-T unshielded twisted pair), should be employed.
2. All client machines connected to the LAN should speak TCP/IP directly.

Item 2 is an important step in establishing a proper foundation for Internet access. Several solutions for Internet connectivity exist that use a server as an AppleTalk-to-TCP/IP gateway or an IPX-to-TCP/IP gateway, reducing the network to utilizing only a single protocol. This approach, however, is strongly discouraged by the IETF's K-12 guidelines. One disadvantage to this approach is that it forces the school to use computers that speak only a proprietary protocol. This "closed system" approach does not guarantee that new architectures and applications based on the TCP/IP protocol suite will work.

The connection of the LAN's in each school building together to form a WAN can be a point of frustration for many school districts. As with LAN's, WAN's consist of the communication medium and the communication protocol. The HorizonNet model uses only the TCP/IP protocol for the WAN, even though AppleTalk and IPX may be on each LAN. Routing of proprietary protocols, such as AppleTalk and IPX, between LAN's should be accomplished by encapsulating the proprietary protocol inside TCP/IP.

4.2 Using the Internet as an Educational WAN

The Internet consists of a collection of LAN's that are connected with various communication links that all speak the TCP/IP protocol suite. The current topology of the Internet, shown in figure 5, consists of various Internet service providers (ISP's) that are interconnected and coordinate the exchange of TCP/IP traffic. Regional ISP's are defined as entities that span large geographic areas with relatively high-speed backbones. Regional ISP's have replaced the large, federally funded backbone networks such as NSFnet. Local ISP's, however, generally service small geographic areas and generally offer more connectivity and support options. All networks connected to regional or local ISP's become part of the world's largest WAN, called the Internet. This interconnection allows people to not only exchange information between each building but to exchange information from any other suitable network that is connected to this highly coordinated worldwide network.

The typical connection to the Internet requires that each school lease a high-speed line from the school building to an ISP. Two distinct considerations are involved in connecting a network to the Internet. The first is the monthly or yearly fee that must be paid to the ISP to route the Internet traffic (TCP/IP) to and from the school. This service charge for routing TCP/IP traffic is sometimes referred to as a membership fee, routing fee, or port fee. The second is the cost of the leased line from the school to the ISP. These communication lines are typically leased from the local phone or cable company. Current pricing for leased lines range from approximately $200 per month for a low-speed connection to several thousand dollars per month for a high-speed connection. Many options are available in selecting the communication line, including Frame Relay, ISDN, SMDS, and T1. All leased-line pricing is expensive in comparison with a standard telephone line. The relationship between the communication line and TCP/IP routing is that the line determines the speed of the WAN connection and the TCP/IP routing gives each machine the functionality associated with the Internet. Both services have certain costs associated with them.

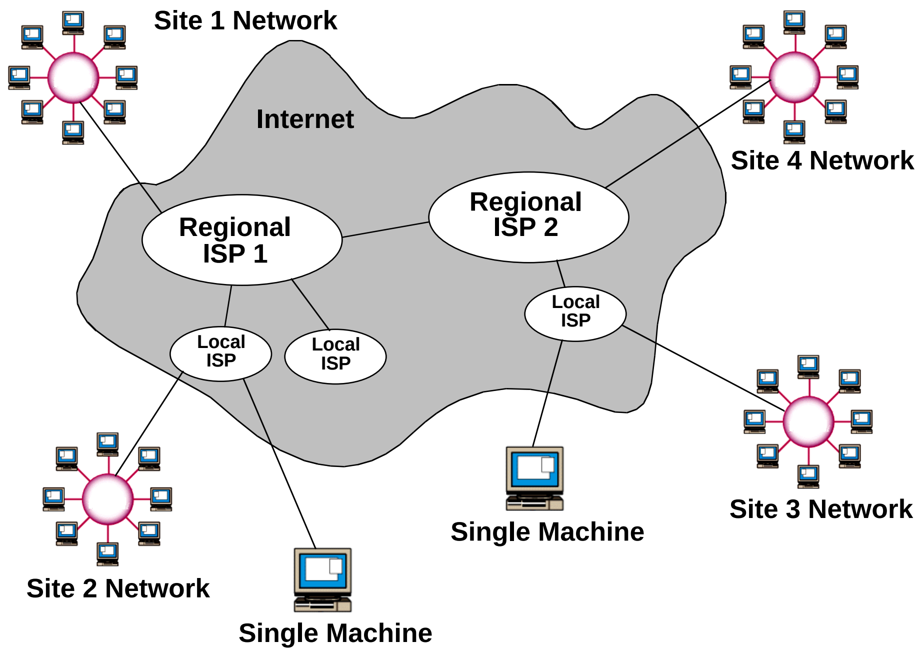

FIGURE 5. Networks and Machines Connected to the Internet

After the communication link and routing of TCP/IP traffic has been established, other services are required for client machines to use TCP/IP. Most ISP's will expect that the proper equipment and expertise for configuring and maintaining services for TCP/IP are located at the connecting site. This includes system administration support for the router and network servers that are located at the connecting site (i.e., the school building in this model). Generally, this task cannot be delegated to a teacher. The California Department of Education K-12 Network Technology Planning Guide states:

> *An organization that implements technology but does not address Network Support as both a technical and personnel issue will soon find itself in trouble. Likewise, organizations that try to glance over network support issues with superficial solutions will be no better off. For example, in many schools, network support issues have been the responsibility of an existing staff member who has been a champion of the technology and who agrees to support the network in*

addition to his/her other responsibilities. While this may be effective in the short term, it has proved to be ineffective in the long term.[2]

The responsibilities of a network administrator are time consuming and require highly specialized skills. Because of the rapid growth of the Internet, the demand for these skills is increasing rapidly. For example, experienced network administrators can easily earn between $40,000 and $60,000 per year in southeastern Virginia.

The recurring costs and required technical expertise associated with leased lines, TCP/IP routing, and network support produce funding requirements that are out of reach for most schools. To offset costs and increase support, the IETF K-12 Internetworking Guidelines suggest connecting individual school buildings to the school district central office and then connecting the central office to the Internet. As a result, the Internet connection can be through the county office, which can presumably help with support. An overview of this model is shown in figure 6.

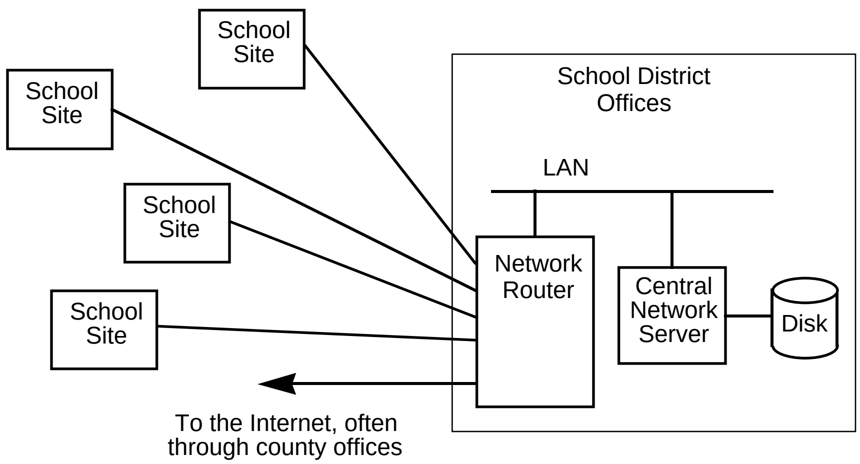

Reprinted from RFC 1709, IETF's K-12 Internetworking Guidelines

FIGURE 6. Interconnection of Schools to the Internet Through Local School District Offices

The HorizonNet model follows the spirit of connecting multiple schools to a central site but goes one step further. Rather than connecting each school building to the school district central office, each school in multiple districts connects to the local K-12 ISP, as shown in figure 7. This K-12 ISP is operated for K-12 schools within the region and consolidates the high-cost items associated with Internet connectivity. In addition to the normal functions of a local ISP, the K-12 ISP customizes the services for the educational community. Examples of these specialized services include network monitoring, content control, and TCP/IP server maintenance.

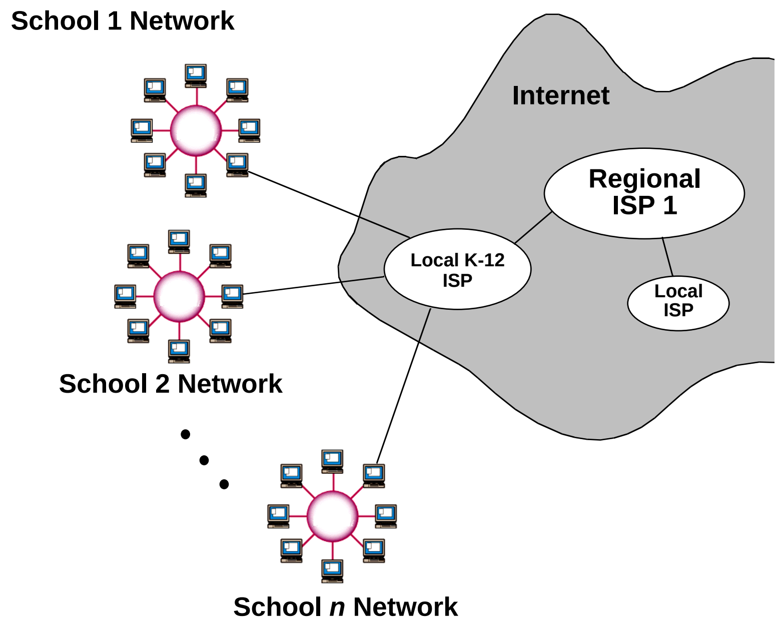

FIGURE 7. Creating a Local K-12 Educational Internet Provider

4.3 Creating the Local K-12 Internet Service Provider

The local K-12 ISP leases a high-speed line to a local or regional ISP and carries the Internet traffic for many school buildings. The K-12 ISP then routes the Internet traffic to the schools by utilizing significantly lower cost methods. In addition, the K-12 ISP can provide a stronger support mechanism for the required equipment in the school than would typically be obtainable from a local or regional ISP directly. Additionally, the system administration support can be customized to meet the unique needs of the lower cost network-to-network connections between the school and the central site. As the need for faster connection grows within the school, the low-speed communication lines can be upgraded to high-speed lines without retraining the users of the system or replacing the computers that define the system.

In this model a low-cost analog telephone line is used to carry traffic to and from a network of computers located in a remote school building. Because the model uses analog telephone lines, which are metered for long distance calls, the local K-12 ISP should be placed in a location in which the cost of leased lines can be minimized and the number of schools in the same local calling area is maximized. Leased lines may need to be used to cross different local calling areas, but the use of these high-cost lines should be minimized. The model presented here assumes that all school buildings are in the same local calling area.

Given that each computer in a school building uses the TCP/IP protocol, IP addresses must be provided for every computer. For the model presented here, every school building receives a class C network address from InterNIC. Each class C network address allows the last three digits (8 bits) of the address to be assigned at the individual school. For example, the network address for Denbigh High School in Newport News, VA, is 204.197.6. The machines at Denbigh High School have addresses within the range of 204.197.6.1 to 204.197.6.254. Because the router also requires an IP address, 253 IP addresses are available for assignment at every connecting site.

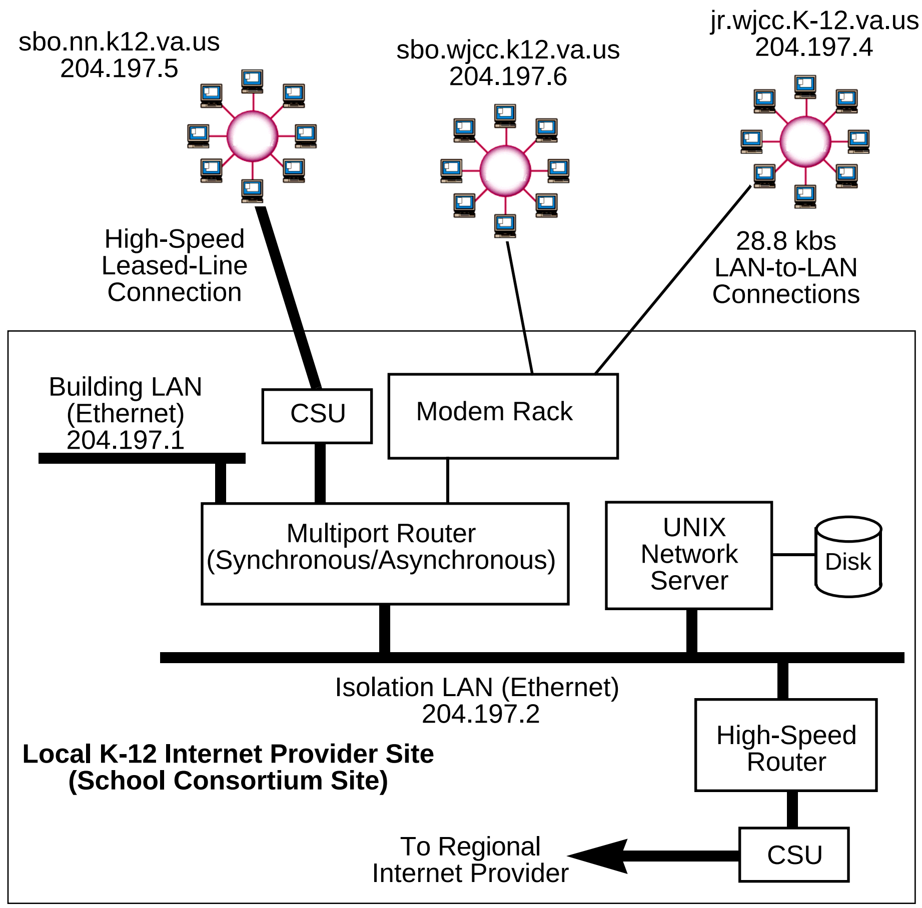

FIGURE 8. Local K-12 Internet Provider Network Diagram

All traffic for the class C network at each remote site is routed from the Internet to the Local K-12 ISP and then to the appropriate school building. As shown in figure 8, an isolation LAN is established at the K-12 ISP for routing TCP/IP traffic through the LAN-to-LAN communication lines. All school building LAN’s must be connected to this isolation LAN with a communication line that is capable of carrying TCP/IP traffic. The bandwidths of these communication lines are determined by each school’s pattern of network use and budget and can range from a standard phone line to a high-speed leased line. A UNIX computer is placed on this isolation LAN to provide primary domain name system (DNS) services for connecting schools. Each school building has a network server for local DNS, but the primary DNS server for all connecting schools is on the isolation network. This configuration is in accordance with IETF recommendations because it allows quick resolution of IP addresses for machines on the global Internet and provides quick address resolution for computers on the remote networks. This prevents DNS services for all computers from going across the slow-speed connections. With the advent of hypertext browsers such as Netscape, proper distribution of DNS services is essential. For any given homepage, a separate DNS lookup is required for each text file, graphic file, and any other file embedded in the page. In addition to providing primary DNS services, the UNIX computer system located on the isolation LAN can be used for other TCP/IP services, such as functioning as the region's news server, anonymous FTP server, and Gopher server.

Many options are available for communication lines that carry LAN-to-LAN TCP/IP traffic and attach to the isolation LAN. A multiport router is connected to the isolation LAN; this router is capable of supporting communication lines that vary from standard analog phone lines, to ISDN lines, to T1 leased lines. Although the multiport router can accept a variety of connection lines, the line with the lowest cost is the standard analog telephone line. However, use of a standard analog telephone line to carry a network's TCP/IP traffic requires that the remote site be designed properly to achieve acceptable results when multiple computers are used. This topic will be discussed further in the next section.

If a standard telephone line is used as a dedicated LAN-to-LAN connection, then the multiport router must dial the remote site and ensure that the line stays open. With this open phone line, point-to-point protocol (PPP) is used to transmit the TCP/IP traffic between the isolation LAN and the remote school building's LAN. If the phone connection is dropped, the router immediately calls the remote site to reestablish the connection. Because each end of the communication line requires a dedicated phone number and because typical charges for a standard phone line are $15 to $30 per month, a dedicated LAN-to-LAN connection can be established for $30 to $60 per month with a performance of 28.8 kilobits per second (kbs). Note, however, that the figure of 28.8 kbs is without compression and also that throughput typically ranges from 30 - 100 kbs with V.42bis compression. Again, this performance can be significantly improved by using distributed TCP/IP servers that run specialized software, as discussed in the following section.

The location of the K-12 ISP (e.g., the New Horizons Regional Education Center for the pilot program) has a separate internal LAN for use by the employees and students. This general-use LAN, although in the same building, is connected to the isolation LAN in a manner that is similar to the connection used for the remote school buildings. In figure 8, this general-use LAN is shown as the building Ethernet LAN that is connected to the multiport router. In effect, the K-12 ISP has two LAN's: a LAN connected directly to the Internet that remote schools connect to and an internal LAN for general use by administrators, teachers, and students at the K-12 ISP site. Configuration of the isolation LAN with the routing for multiple networks and supporting LAN-to-LAN communication links in the range of $30-60 per month provides the remote school buildings with an affordable recurring cost for dedicated Internet access for the entire LAN. As a demand for higher bandwidth is demonstrated, the isolation network can accept higher speed leased lines.

To create the isolation LAN, the following items must be obtained:

- Through a regional ISP and InterNIC, obtain routing for multiple class C networks as determined by the number of schools that will eventually connect. This cost is recurring and typically ranges from $20K-30K per year, depending on the level of support.
- Obtain a high-speed leased line (T1 or better) to the ISP. This line is normally leased from communication companies such as Bell Atlantic, MCI, or even the local cable company.
- Obtain all hardware, software, and phone lines for network-to-network connections at the center, as detailed in the appendix.

4.4 Connecting Remote School Buildings

The connection of each school building's LAN to the isolation LAN of the K-12 ISP is discussed in this section. As mentioned previously, the costs associated with building a communication infrastructure that is highly scalable can put the cost of LAN-to-LAN connectivity out of reach for most school buildings. School buildings, with classroom and administrative activity, represent relatively large networks. Given the potentially large number of computers per school building and the network infrastructure required, the IETF K-12 internetworking guidelines must be followed to ensure the scalability of the network design as demand increases. The infrastructure model presented here is designed for scalability by using open systems and freely available software.

Two basic models are available for connecting to the Internet. The first model, depicted in figure 9, is to deploy high-speed leased lines to all school buildings and centralize the support at the school board office. All TCP/IP services can then be performed at the central site. Because most or all server functions are performed at a central site, each remote LAN requires a high-speed communication line so that each client machine realizes high-speed access. Additionally, the centralized TCP/IP servers must be able to handle the large load presented by thousands of machines that are using the Internet simultaneously.

A simple example of how a centralized server approach for an entire school division can become overloaded is given with just the delivery of electronic mail. Consider a school division with 30 school buildings and an average of 200 computers per building on the Internet. If the computers constantly run a post-office protocol (POP) mail client and check for e-mail every minute, the central e-mail server will be required to support 360,000 requests per hour for the district. Furthermore, the bandwidth of the communication link will have be large enough to support 200 users continuously checking for e-mail on the central server, even though most may not receive any mail during these requests. A single user running a POP mail client, such as Eudora, will typically log in and check for mail (automatically) about 360 times during the day. For a user that receives 10 mail messages per day, less than 3 percent of the server log-ins will result in mail delivery. The remaining 97 percent of server log-ins will produce no e-mail delivery but will compete with other users for the WAN bandwidth use. This simple example shows how quickly a centralized service, such as e-mail, can overload this design.

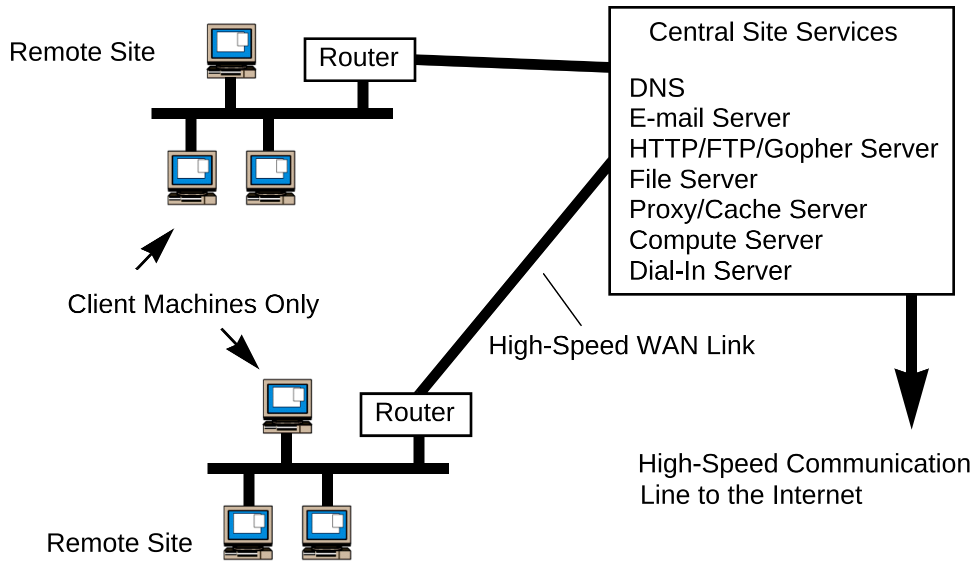

FIGURE 9. Centralized Model for all TCP/IP Servers

The approach recommended in this paper provides high-speed access to each user at the remote site by distributing the TCP/IP server functions to each building. Access to this TCP/IP network server is then made over the LAN, which is generally much faster than the WAN communication line. For example, users will send and receive e-mail at 10 megabits per second (Mbs) and not the lower WAN speed. After a user sends e-mail to the server, the server releases the user's computer and then forwards this file to the Internet over the slower WAN communication link. From the user's perspective, however, the e-mail was delivered at the faster LAN speed of 10 Mbs. This local server, which is a high-performance UNIX-based computer, is leveraged to support other functions associated with the TCP/IP protocol.

The distribution of small, powerful UNIX-based systems is more efficient than a large, centralized computer system that is capable of large loads. The computer industry now produces small workstations capable of handling small to medium loads at a much lower cost than large mainframes. A well-designed distributed-server topology can cost less than an equivalent centralized-server architecture. Additionally, a distributed-server approach allows the remote sites to reduce the required bandwidth of the WAN communication link and save on this major recurring cost item.

4.5 Scalability

Each school is responsible for providing and paying for the communication line located at the school. At a minimum, this line is simply an analog telephone line. As demand for network speed increases, several options are available for upgrading the communication line to the connection hub. Many third-party vendors manufacture expansion cards for the Sun workstation that support ISDN, Frame-Relay, and dedicated T1 leased lines. The type of communication line used does not effect routing or server functions. The communication line is a simple upgrade and does not require rebuilding the entire communication infrastructure.

5.0 Maximized Performance Using Innovative Technologies

Although most local and wide area networking models deal with obtaining the highest capacity communication lines possible, one of the largest expenses associated with most architectures is the recurring cost of these WAN communication lines. By analyzing the typical usage patterns of the various TCP/IP services, technological and operational strategies may be implemented that maximize the use of existing and readily available low-cost communication lines.

5.1 Characterizing TCP/IP Traffic Patterns

Given the TCP/IP services required to run popular Internet activities such as e-mail, WWW browsing, and voice/video delivery, the TCP/IP traffic that traverses the WAN link can be categorized into four areas for a distributed server approach. (see fig. 10.) These traffic patterns, categorized from most efficient to least efficient are: store-and-forward; network-based caching; split or broadcast; and direct WAN access.

With a local network server, the bandwidth required by a given LAN can be greatly decreased by using software that will take advantage of alternate data retrieval schemes (i.e., other than direct WAN access). The previous example with e-mail is an example of a store-and-forward traffic pattern. This approach provides all users on the LAN with fast access to this service, regardless of the WAN communication speed. As shown in figure 11, the model presented here uses software on the local UNIX network server to minimize direct access to the WAN communication link. Many Internet activities, such as WWW browsing and e-mail, can use data retrieval schemes other than continuous direct WAN access. In effect, the model detailed here provides client machines that are running Internet software such as Netscape with “virtual bandwidth” to the Internet.

The network services that are part of a typical Internet connection model for an efficient, highly utilized network are shown in figure 12. Some of these functions could be placed outside the LAN at a central location; however, those functions would require repeated connections using the slower WAN communication line. Separate computer systems can also be used for each of these functions. The model presented here combines all functions shown in figure 12 into a single high-performance UNIX computer system;

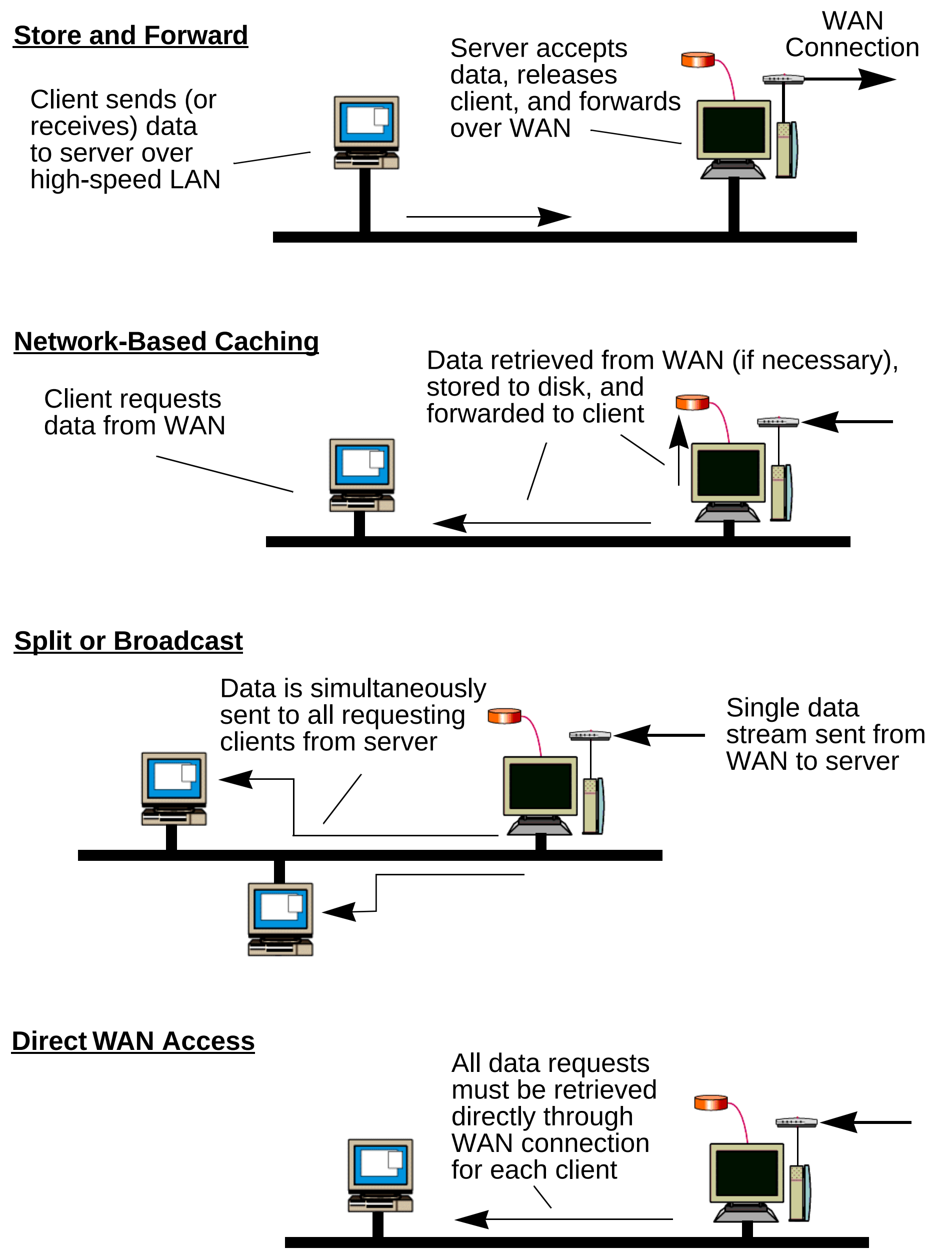

FIGURE 10. Four Basic Traffic Patterns for Local Network Server

these functions are placed directly on the high-speed LAN so that the performance at the desktop level is much higher than that which would be suggested by the speed of the WAN connection. The software required to perform these server and router functions is available free from the Internet for K-12 institutions or is included with the UNIX system

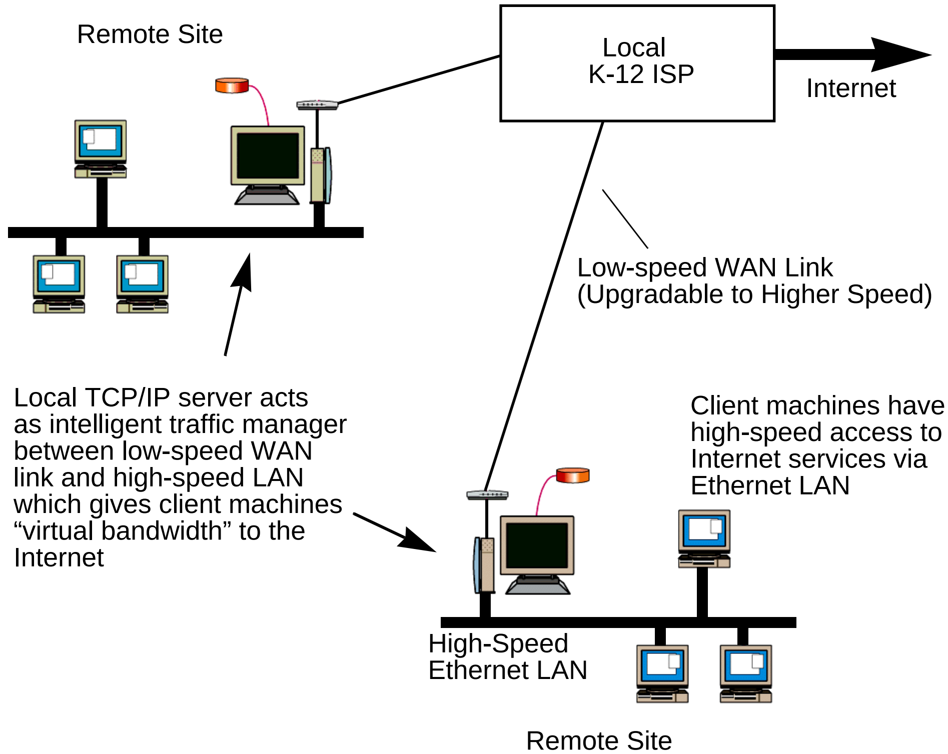

FIGURE 11. Distributed Server Topology

5.2 Efficient Use of Existing Bandwidth

Understanding the various data-transmission schemes outlined in the previous section is important; appropriate implementation of new and existing technologies can drastically improve network performance from a user's perspective at any desktop computer on the LAN. The following functions are critical elements in this analysis.

5.2.1 Routing

Any internetworking model requires the use of routers to direct the LAN and WAN traffic. Routing functions can be handled by a dedicated hardware device that is located between the LAN and WAN communication lines. These dedicated routers typically cost between $1000 for a low-speed connection and $3000 for a high-speed T1 connection. In addition to the router, a CSU/DSU is required, which acts as a specialized modem for the WAN communication line.

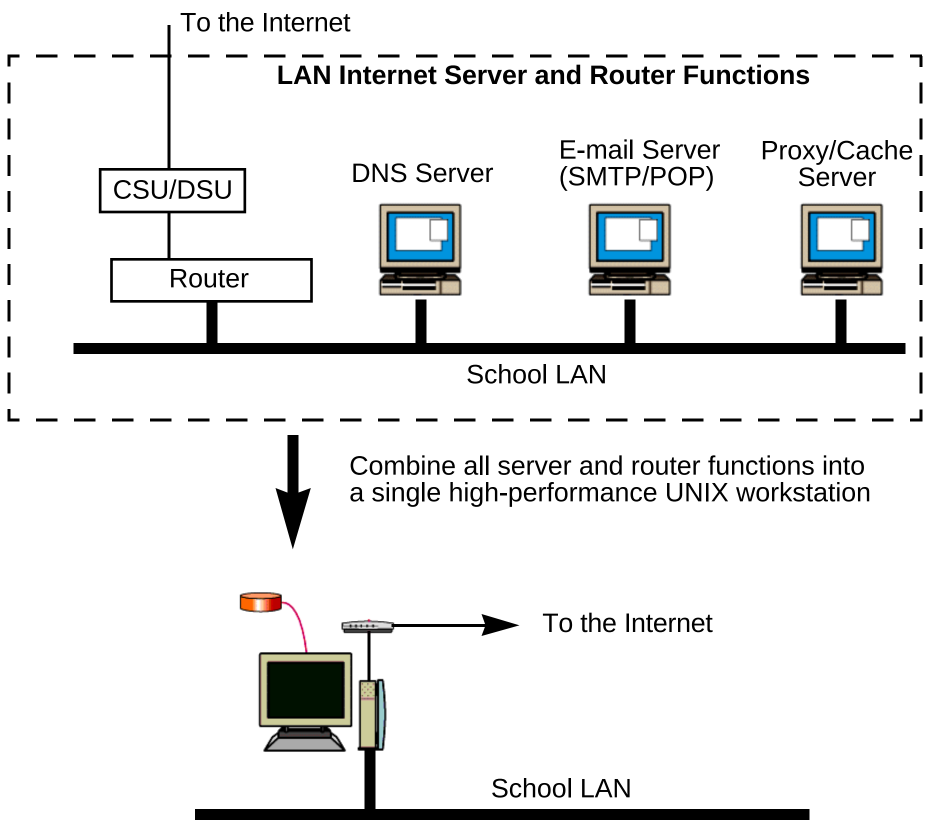

FIGURE 12. All Local Router and Server Functions Combined into a Single UNIX Computer System

The model presented here uses software that runs on the UNIX computer system and a specialized expansion card for the CSU/DSU connection. The lowest cost connection is a simple high-speed 28.8 kbs modem in conjunction with a high-speed serial card. For this low-cost LAN-to-LAN connection pipe, performance can range from 30 - 100 kbs, depending on the type of data to be transferred. Additionally, expansion cards are available allowing upgrades to ISDN, Frame Relay, T1, and ATM communication lines.

The TCP/IP protocol is transported over a standard analog telephone line by using PPP software. The software for PPP in this model is dp-2.3, which is available from Purdue University.

5.2.2 Proxy/Cache Service

The Internet connection model in which all or most TCP/IP services are centralized creates a situation in which every data transfer request must go out over the slower WAN link and retrieve data from a remote server. The advent of the WWW has produced a situation in which many graphics, sounds, and video clips are sent from remote servers to client machines. Problems can arise with connection speeds when many users are trying to access outside data at the same time. For example, in a training session students are told to connect to www.whitehouse.gov. Each computer opens a separate connection and downloads the data associated with this site, as depicted in figure 13. If 20 computers are located in the classroom and the leased WAN communication line is 56 kbs, users will each experience transfer rates equivalent to a 2400-baud modem.

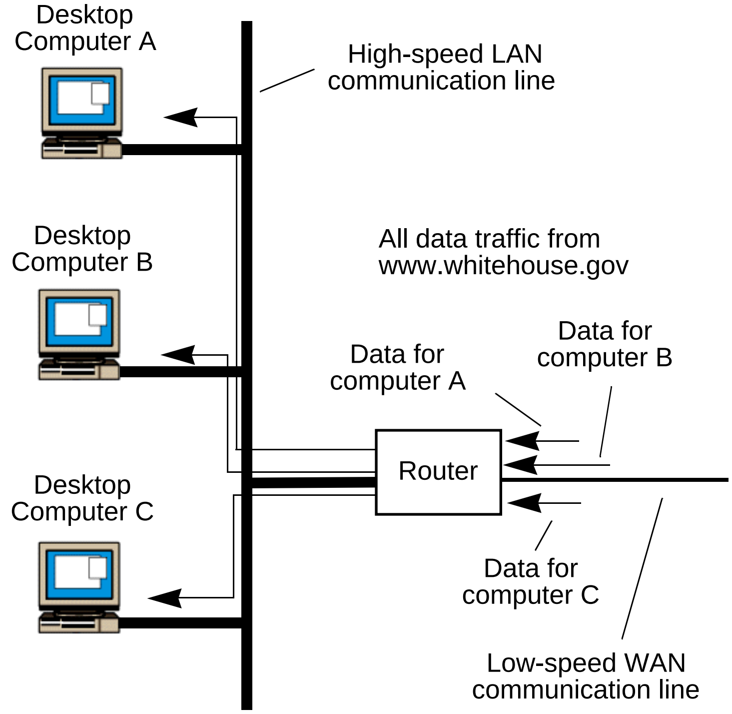

FIGURE 13. Separate Connections to Remote Server for Individual Client Machines

The traffic pattern described in the previous scenario can be used to design a system in which the duplicate traffic is removed from the WAN communication line. Rather than require each computer to connect separately to the outside server, a proxy/cache server is used to remove the repetitive traffic from the WAN communication line. Instead of requiring each desktop computer to contact the remote host directly, each computer contacts a networked proxy/cache server, which in turn connects to the remote computer. If multiple connections are requested to the same remote site, the proxy/cache server only downloads the data once and simultaneously forwards that data to the requesting computers. The proxy/cache server also caches the data on a large disk. As data are downloaded from a particular site on the Internet, the data are simultaneously stored on a large hard disk and forwarded to the requesting computer. (See fig. 14.) Any future request for data that have been previously fetched and stored on the proxy/cache server will realize a much faster throughput. (See fig. 15.) The speed at which a user receives requests from the WWW is now dictated by the LAN speed and the network server speed, assuming that the information requested has not changed. For items that have been previously cached, the proxy/cache server verifies that the data have not been updated by exchanging approximately 50 characters with the remote site. If the data have not been changed, the requesting computer receives the cached data. The proxy/cache server runs on the same UNIX-based computer that is used as the router.

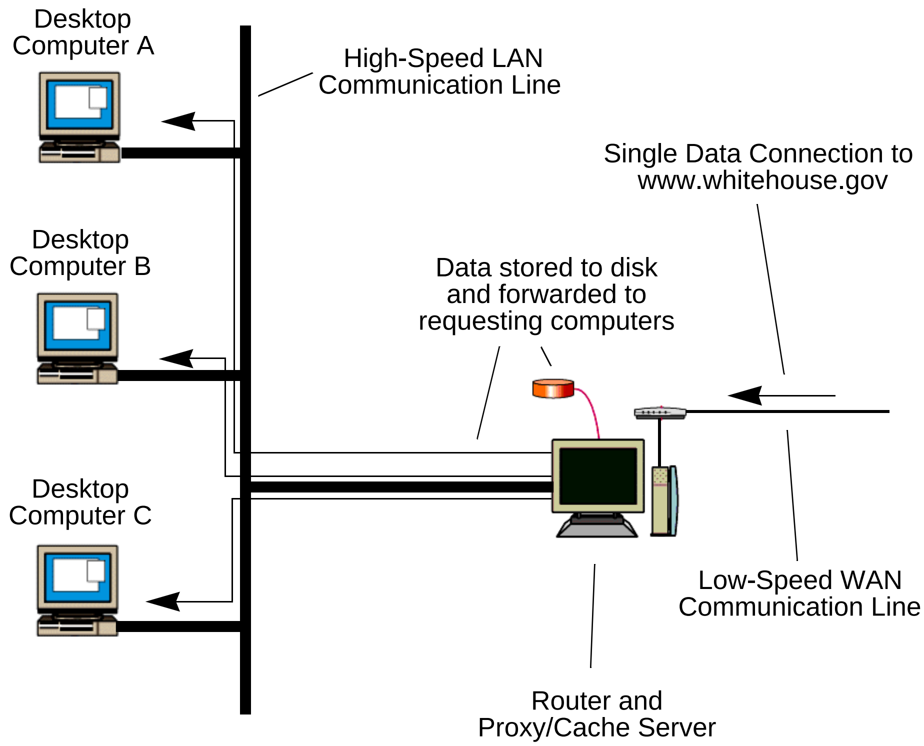

FIGURE 14. Operation of Proxy/Cache Server

The use of a proxy/cache server can greatly reduce the data volume from remote Internet servers for repetitive network traffic patterns when the accessed data are fairly static. Caching of frequently accessed data benefits not only the client computers on the LAN; the remote Internet servers also realize less load because they do not have to deliver data to each client machine on the Internet.

The amount of disk space set aside for cached data is important for the performance of the proxy/cache server. Although no optimal size of the disk partition for cache data has been determined, approximately 1 Gbyte of disk storage is currently set aside for cached data. Typically, the operational cache size is approximately 500 - 700 Mbytes, and cached data are deleted if not accessed in 6 weeks. Obviously, the algorithms used to determine when to remove cached data depend on several variables, such as the WAN communication line speed, the size of the file, and frequency of access. This topic would benefit from further research.

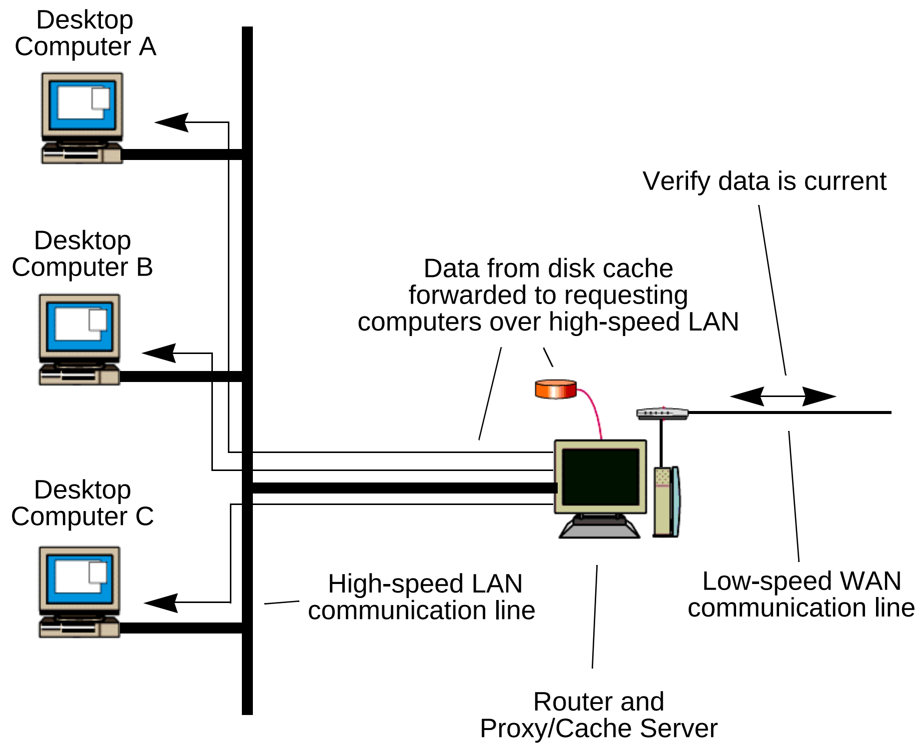

FIGURE 15. Obtaining a Cache Hit on the Proxy/Cache Server

Another major benefit to running a proxy/cache server is the issue of content control. School systems must be able to control, to some degree, what is accessed via the Internet. Because the proxy/cache server checks each uniform resource locator (URL) against its cache database, access privileges can be established for any given URL. Specifically, lists can be set up that restrict or allow access to specific documents. Although access to specific domains and IP addresses can be blocked or allowed at the router level, this task is better performed at the URL level. Although some domains, such as playboy.com, should obviously be blocked, instances occur in which some information on a particular server is valuable to education and other information on the same server is objectionable. The Yahoo Internet Directory is an example; this directory lists thousands of Internet servers by category.

Several proxy/cache software packages are available from the Internet; two of these have been tested in the pilot program. The first server, and one of the most popular, is the CERN httpd server. In addition to performing the functions of a WWW server, this software can also cache HTTP, FTP, and Gopher traffic. Several problems were encountered with this software package in its original form. First, the CERN httpd server would not stop downloading large files even if the requesting client computer stopped the download. For example, if a user clicked on a WWW home page that contained a 300-

kbyte image, clicked the stop button, and then clicked on another page with a 400-kbyte image, the CERN httpd server would continue to download both images even though the user discontinued the 300-kbyte file. This problem in the initial testing of this software created situations in which large files that were not going to be used were competing for bandwidth. The second problem that occurred with the CERN httpd server was in the algorithm used to determine when files were to be cached or discarded. The CERN algorithm placed greater value on disk space than on the bandwidth or the time required to download the file. For low-speed connections, the time needed to download a file is more valuable than the disk space required on the server.

The proxy/cache software developed at the NASA Langley Research Center for this pilot program is based on the Ichthus cache server software written by James Mathew Farrow at the University of Sydney, Australia. The original Ichthus proxy/cache software is written in the Perl language and is still immature in development. Although fairly stable, the software only caches HTTP traffic. The author of Ichthus continues to work on this software as time permits. The Ichthus proxy/cache server does not share the performance problems of the initial CERN httpd server and, because of its simplicity, is fairly easy to modify. This Perl-based proxy/cache software has been modified to include a filtering mechanism for URL’s, which was discussed previously. In addition, this software is currently being rewritten to include functionality that will increase the cache hit rates achieved on the school LAN. The rewrites of this software are discussed later in this paper.

The performance of the proxy/cache server was demonstrated at the Yorktown Elementary School during the spring of 1995. A Sun Microsystems SPARCstation 5 was placed in a school computer laboratory that contained 31 Apple Macintosh computers. The computers were networked with thin Ethernet cabling and were given static IP addresses. The SPARCstation 5 was connected to the Internet with a standard analog phone line and a 28.8 kbs modem. All students in the elementary school were cycled through this laboratory to learn about and use the Internet with the Netscape WWW software.

The HTTP cache hit and fault rates for the LAN during the first 5 days of Internet use are shown in figure 16. A cache hit occurs when a client machine requests information with a WWW browser and the current information is already stored on the server. A cache hit results in delivery of the data at the high LAN speed. A cache fault is when the proxy/cache server does not have the requested information stored or when the information stored locally has changed on the remote server (verified by an if-modified-since request).

The activities represented in figure 16 represent focused Internet activities. The first day’s activities included the initial setup and exploration by the teachers of the school on a server that already contained 500 Mbytes of cached data. The second and third day’s activities included bringing every student in the school through the computer laboratory for approximately 20 minutes and taking them on a virtual tour of the White House, with the primary goal of teaching students how to navigate the Internet with the Netscape software. During days 2 and 3, approximately 250 Mbytes of HTTP data were delivered to the client machines, with approximately 17,000 requests per day made to the proxy/cache server. During the fourth day and later, the students were given specific projects to work on. Even activities that use the Internet freely and utilize data that are probably not already cached can produce high cache-hit rates. For example, many students were instructed to search for information on endangered species. This assignment produced a situation in which most students on the 31 client machines typed “endangered species” into one of the WWW search engines. Search-engine requests are not cached; however, they will return similar hypertext links, so students still visited many of the same sites other students had already downloaded or were downloading.

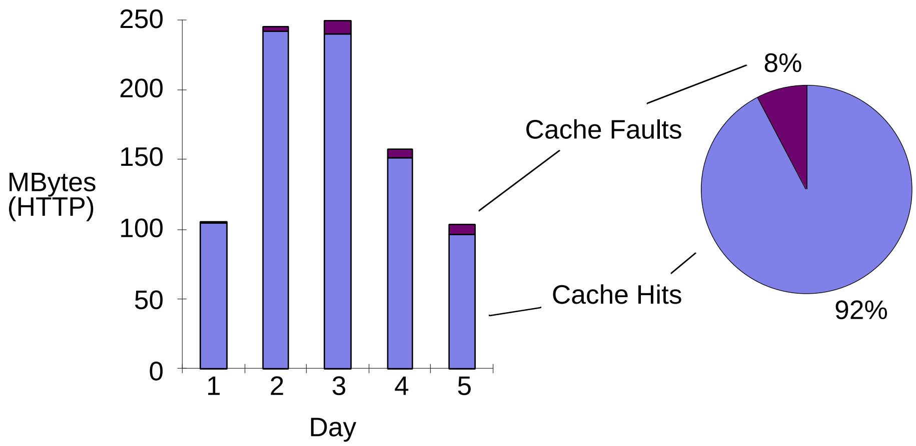

FIGURE 16. Cache Hit and Fault Performance of Yorktown Elementary School Network

The cache hit rate shown in figure 16 can be achieved for focused activities, as demonstrated in these pilot studies. However, unfocused Internet “surfing” by client machines on the LAN will decrease the cache hit rates. Realistic performance figures indicate that cache hit rates can range from 80-95 percent for focused activities and from 30-50 percent for unfocused activities. These performance figures assume the proxy/cache server has substantial previous use and is not a new or seldom-used service.

One example of current activity at NASA Langley Research Center is the software rewrite of the existing cache server. The modified version currently includes a filtering mechanism; however, the software is being rewritten to precache a series of URL’s during times when the WAN connection is not in use. For example, when CNN’s homepage is updated, the proxy/cache server will automatically update the local cache file before a request comes in from a client machine. The proxy/cache server is also undergoing modification to meet the needs of K-12 schools in verifying users on client machines and monitoring all incoming and outgoing traffic. This software development is based on open standards so as not to favor any particular type of client machine. The motivation for this development effort is primarily to research algorithms for minimizing direct WAN access over low-speed links, rather than to compete with other commercial proxy/cache software products.

5.2.3 DNS Service

Another important function for the local TCP/IP server is to resolve host name addresses to the numerical IP address. This resolution is referred to as DNS service, and this translation must occur for every e-mail message sent and every Netscape hypertext link that is selected by a user. For the model presented here, DNS is performed for the local network on the local UNIX-based computer. The local DNS server always knows how to resolve local machine names, so that local TCP/IP services such as local e-mail do not have to depend on name resolution occurring over the low-speed link. Additionally, all DNS requests that must traverse the low-speed link are cached on the UNIX server so additional requests for the same site are handled locally. The primary name service for the network, however, is performed by the local K-12 ISP so that requests from the outside Internet do not have to traverse the low-speed WAN link to resolve names at the remote sites.

5.2.4 Electronic Mail

Creation of a local e-mail server also significantly increases the apparent bandwidth to the Internet. The UNIX system uses a combination of the standard UNIX sendmail program in conjunction with a POP mail server. This combination allows users to use client mail packages, such as Eudora, which run on common desktop platforms. E-mail is sent and delivered to the Internet over the low-speed WAN link, but the user exchanges mail with the UNIX server at the higher LAN speed. Users can then exchange large files over the Internet at the apparent bandwidth of the LAN.

Table 2 categorizes how the various networking services are implemented in the architecture described in this paper.

TABLE 2. Function and Server Software for Dividing the TCP/IP Traffic Patterns

Data Transfer Method	Function	Software on Server
Store-and-Forward	e-mail	sendmail
	DNS (local clients)	IchthusK-12
	HTTP (precached)	IchthusK-12
	FTP (precached)	IchthusK-12
Network-Based Caching	HTTP	IchthusK-12
	FTP	IchthusK-12
	DNS	named
	Gopher	IchthusK-12
Split or Broadcast	MBone	None
	One-to-many video	CU-SeeMe
	One-to-many audio	RealAudio
Direct WAN Access	Point-to-point video	None
	or audio conferencing	
	HTTP (CGI-bin applications)	

6.0 Leveraging the Internet Server

The purchase of a UNIX computer system for each school building represents a large capital expense for any school district. This capital cost can be justified if the computer system is capable of performing tasks that would otherwise have to be performed by additional computers. The UNIX operating system is ideally suited for handling multiple tasks without degrading overall system performance. From its inception, UNIX was designed for a networked environment, whereas for many of the operating systems found on today's desktop computers networking capabilities are being retrofitted into what was once a stand-alone system. More than two decades of development of UNIX-based networking services by universities, government laboratories, and industry have lead to an operating system that is extremely stable, capable of supporting multiple simultaneous users, concurrent processes, and heavy input and output demands. By taking advantage of this powerful and stable operating system, schools can avoid the expense of multiple servers by utilizing the inherent power of a UNIX platform to provide many common networking services. (See fig. 17.)

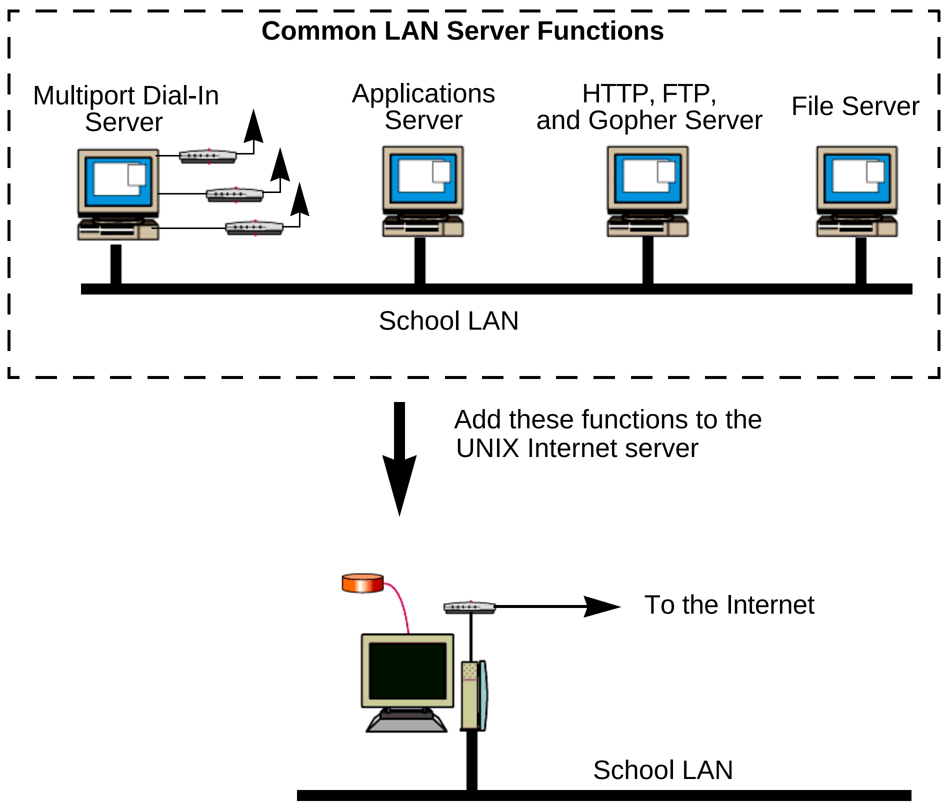

FIGURE 17. Combining Common Server Functions by Using the UNIX Platform

6.1 File Server

The UNIX server can be used to provide file service for the client computers. Because the UNIX operating system provides a true multiuser environment, every user can be given a secure space on the system to store files. The problem that occurs in using the UNIX machine as a file server is that most client computers in schools use the AppleTalk or Novell IPX protocol for file service. The cleanest approach would be to require the client machines to use a network file system (NFS) software product; however, most schools already have some infrastructure built around AppleTalk or IPX.

The UNIX systems in use for the HorizonNet program use the Columbia AppleTalk Package (CAP) software, which allows the Sun workstation to speak native AppleTalk, in addition to the standard TCP/IP. Users on the network can mount their home directory on the UNIX workstation as a standard AppleShare volume by simply using the Chooser and entering their UNIX login name and password

Software products are currently available that allow UNIX computer systems to speak IPX and, thus, become a Novell file and applications server; however, these products have not yet been tested fully in the pilot program.

6.2 Dial-in Server

The UNIX system can also provide dial-in service to the school's LAN and, hence, the Internet. By using a multiport serial card, additional modems and phone lines can easily be attached. Each school building can then allow administrators, teachers, and students to dial in by using a machine-to-network PPP connection. This approach to dial-in service is much easier to scale because the load of assigning user accounts is distributed to each building.

The software for creating a dial-in server is the same dp-2.3 PPP software used to maintain the communication link between the remote site and the connection hub. The software is used in conjunction with a multiport sbus serial card that is capable of full modem control on each port and can transmit data at 115,200 bps on each port.

6.3 HTTP, FTP, and Gopher Server

The UNIX system can also be used as the primary WWW server for each school building by using the httpd software written by NCSA; as a WWW server, the UNIX system would be named www.*school*.*district*.K-12.*state*.us. Although the UNIX system can also be used as an anonymous FTP and Gopher server, these services have not yet been implemented in the initial testbed.

6.4 Computational Server

In addition to the communication infrastructure, utilization of the network for executing applications is also important. The computational infrastructure must mesh with the communication infrastructure for any model to be successful. The communication infrastructure previously outlined can support efforts that require computational activity in a client-server relationship. These efforts include:

- Mathematical analysis and simulation software
- Library automation software
- Programming software
- Word-processing software

These applications can physically run on the local UNIX workstation; the display is piped back to the client machines on the network by using X-window software. This configuration allows schools to achieve higher performance on older computers that are capable of running an X-window software package but not high-end scientific software.

6.5 Summary of Required Software and Hardware

In the model presented here, the connecting school building is responsible for installing all networking and computer systems that are capable of speaking the TCP/IP protocol. Additionally, the school provides an Ethernet connection port for the UNIX workstation and all necessary communication lines. Figure 18, which is a composite of figures 12 and 17, provides a graphical representation of each function on the remote LAN for which the UNIX workstation is responsible.Tables 3 and 4 describe the configuration of the hardware and software for the UNIX workstation at each school building.

TABLE 3. Basic UNIX Workstation Breakdown

Equipment	Description	Approximate Price
Sun Workstation	Sun SPARCstation model 5, 85 MHz CPU, 1 Gbyte internal disk drive, 32 Mbytes RAM, 17 in. color monitor, internal 4x CD-ROM drive, internal floppy disk drive.	$5000
2.1 Gbyte Disk Drive	Internal disk drive for Sun SPARCstation model 5.	$900
DAT Tape Drive	4 mm DAT tape drive with SCSI-2 interface for use with SPARCstation model 5.	$900
Multiport Magma Serial Card	High-speed serial card capable of full modem control and 115,000 bps data transfer. 2Sp+1 card supports two high-speed serial connections. The 4 Sp can be used for additional serial connections.	$350
Telebit FastBlazer Modem	Modem capable of 28.8 kbs and V.42bis compression. Software upgradable.	$750

TABLE 4. Software Suite for UNIX System

Function	Software	Available from
Router	Standard UNIX **routed** can be used although not required because each remote site has only one default route	Included with Sun workstation.
Proxy/Cache Server	**Ichthus-K-12** is a modified version of the original **Ichthus** cache server	NASA Langley Research Center.
DNS Server	Standard UNIX **named**	Included with Sun workstation.
E-mail	Standard UNIX **sendmail** is used in conjunction with a POP3 server called **popper**.	**sendmail** is included with the Sun workstation and **popper** is available from ftp.qualcom.com.
File Server	**CAP** (Columbia AppleTalk Package) is used for AppleTalk-based file service. **UAR** (UNIX AppleTalk Router) is used to establish AppleTalk zones	Columbia University
Dial-In PPP Server	**dp-2.3**	Purdue University

6.6 Determining WAN Bandwidth Requirements

The number of computers that can be supported by this network design is much greater than most would expect. The class C network addressing that is used in this design allows 253 computers on the LAN to have unique IP addresses. Although this network design can easily support 253 computers, the number of computers that simultaneously use the Internet will determine the bandwidth of the WAN communication line required to achieve acceptable performance. For example, if a school has 100 computers with IP addresses connected to the LAN and 3 are running Netscape, 4 are running TurboGopher, and 50 are running Eudora (e-mail), only 7 computers would be simultaneously running interactive Internet applications. Electronic mail is not typically counted in this number because its required bandwidth is relatively small and because e-mail uses the store-and-forward method to transfer data.

The topology of the LAN also must be considered in determining the required bandwidth for the WAN. A school with large computer laboratories in which students are all accessing the Internet simultaneously will have different WAN bandwidth requirements than a school with three or four computers in every classroom. (A computer laboratory in which a large percentage of the class is doing a similar activity is an example of synchronous activity, whereas a distributed network of three or four computers in each classroom is an example of asynchronous activity.) Additionally, how the Internet is used greatly affects the required bandwidth. Focused activities produce cache hit rates in excess of 80 percent on the local UNIX server; unfocused exploration can bring the cache hit rates down to 30-50 percent.

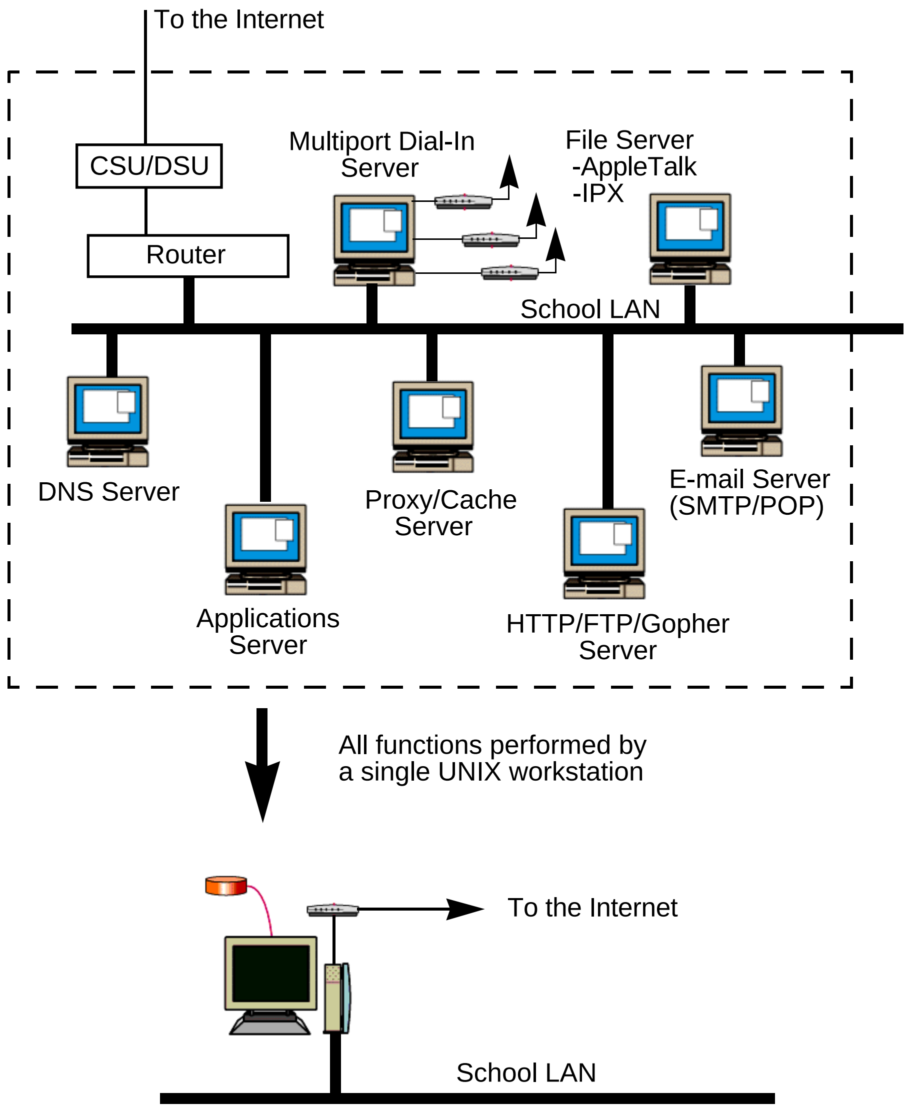

FIGURE 18. Services Provided by the UNIX Workstation

7.0 Distributed Technical Support and Training

The technical support of a computer and network infrastructure can be one of the most costly items. For example, UNIX systems have a reputation of being difficult to configure and maintain. However, the UNIX systems and the software suite used in this pilot program, once configured, were extremely stable and have required only periodic maintenance and troubleshooting. To minimize the cost and maximize the effectiveness of a support structure, a two-tier model for supporting the UNIX systems was developed during the pilot program.

The two-tier technical training and support model presented here requires the local K-12 ISP to provide the following:

- Primary system administrator (approximately $40K-60K per year).
- Training and support for secondary system administrator in each school.
- Training classes for teachers and administrators on using the Internet.

7.1 Primary System Administrator

The connection hub will be required to provide a primary system administrator that will be shared by all connecting schools. This system administrator should be responsible for the primary care of the single UNIX computer in each connecting school, as well as all equipment on the connection hub Isolation LAN. The primary system administrator should have the following responsibilities.

7.1.1 Internet (TCP/IP) Routing.

The primary system administrator should be responsible for ensuring that the routing hardware and software are working properly. This responsibility includes both LAN's at the K-12 ISP and the low-cost connections to the schools.

The primary system administrator should also be responsible for the primary nameserver at the K-12 ISP and for adding and deleting network addresses. Therefore, all equipment that uses the TCP/IP network at the school must be coordinated through the primary system administrator for the purpose of assigning network TCP/IP Addresses.

As the need arises, the primary system administrator should be responsible for evaluating and upgrading the system to provide faster access for each school. This activity is coordinated with each school district and the local K-12 ISP.

7.1.2 Global System Security

The primary system administrator should be responsible for evaluating and implementing appropriate system security measures on a global basis. For example, if firewall software becomes stable and necessary, the primary system administrator should be responsible for implementing it. Also, the primary system administrator should be responsible for installing security patches and upgrades for the software on the primary UNIX machine at each school.

The primary system administrator should also be responsible for ensuring that reasonable security procedures are implemented at each school. This task should be accomplished primarily through coordination with the secondary system administrators and the appropriate use of security programs such as Crack.

7.1.3 Training

The primary system administrator should be responsible for training and, optionally, certifying each secondary system administrator selected for the school. In addition, a “cookbook” should be provided to each secondary system administrator to detail some common problems and solutions. This guide is intended to be a working document that is modified as necessary in coordination with the secondary system administrators.The primary administrator may also want to establish a listserv to which all of the secondary system administrators belong. A listserv allows new secondary administrators to benefit from the knowledge of experienced colleagues, provides a forum to identify items that can be collated into a FAQ, reduces the burden on the primary system administrator, and allows the primary system administrator to quickly identify common problems that may need to be addressed in a timely manner.

7.1.4 UNIX System Configuration

As each school is attached to the network, the primary system administrator should be responsible for the configuration of the UNIX machine that provides the school's access to the K-12 ISP network. The primary system administrator should also be responsible for updating the hardware and software for the primary UNIX machine at each school.

The primary system administrator should also be responsible, with the help of the secondary system administrators, for ensuring that the primary UNIX machine at each site is up to date with the other sites (i.e., one system shouldn’t be at a much higher software and OS revision level than the others).

7.1.5 Software Utilities

The primary system administrator should be responsible for providing scripts as necessary to ease the burden on both the primary system administrator and secondary system administrators. These scripts and programs should be done in coordination with the secondary system administrators as they are needed.

7.1.6 Hardware and Software Evaluation

The primary system administrator should be responsible for the evaluation and implementation of any hardware and software to improve the access throughput of each system.

7.2 Secondary UNIX System Administrator

Each school that contains a network should provide at least one person to perform the secondary system administration for the UNIX workstation located at the school. This person will be the primary point of contact for everyone in the school building in regard to issues that concern the network connection, as well as services provided by the UNIX workstation. The group of people responsible for secondary system administration at all schools should meet on a periodic basis with the primary administrator (possibly electronically) for discussion of system administration issues.

This function is extremely important and must be handled professionally. If passwords are not set correctly, the data on the UNIX computer could be jeopardized, which in turn could jeopardize the Internet connection. The administrator must understand the issues related to computer security and ethics.

All secondary system administrators should complete a training program. At a minimum, the secondary administrator must receive training and have a working knowledge of the following:

- Basic UNIX
- Basic networking
- Computer security and ethics

This training should be provided by the local K-12 ISP. Training materials for these courses are available from NASA Langley Research Center and were developed for the pilot HPCC K-12 Program. The duties of the secondary administrator at the school are outlined below.

7.2.1 Account Maintenance on the Local UNIX Server

User accounts that are required on the local UNIX workstation should be added and maintained by the secondary system administrator. These accounts are required for certain Internet activities, (e.g., e-mail) and for any of the functions associated with storing files on the UNIX machine (e.g., Apple File Sharing).

7.2.2 Acceptable Use Policy Training

A brief training period for all teachers and students should be provided during the year for individuals who want to use the services available on the Internet. Teachers and students should be made to understand and abide by the acceptable use policy (AUP) established by the school for both the computers and the network. Because millions of people on the Internet can attempt to gain access to the computers on the school network, students and teachers who use the computers must understand the importance of establishing proper passwords. Lecture notes for a security and ethics class along with AUP's from various schools are available from NASA Langley Research Center.

7.2.3 Automated Procedures

The secondary system administrator should assist the primary system administrator in monitoring the automated functions of the system (e.g., ensuring tape backups are made and notifying users if their password is insecure).

7.2.4 Coordination of Network Expansions with Primary System Administrator

The primary system administrator must be notified of any computer, router, or gateway that is connected to the Ethernet network. All devices connected to the network that use the TCP/IP protocol must be assigned an IP address and a valid network name. This name and IP address are programmed into the primary name server, located at the local K-12 ISP, and the UNIX system located at the school. The secondary system administrator should keep a record (spreadsheet) and coordinate with the primary system administrator on the assignment of all IP addresses and names located on the school network.

7.2.5 Training

The secondary system administrators receive ongoing training from the primary system administrator on new procedures and policies regarding the computers and networks. Training time should be allocated to the secondary administrators for this purpose. When applicable, secondary system administrators should then train appropriate users at their school.

8.0 Conclusion

The key to successfully providing low-cost Internet connectivity is to minimize the number of times that client machines on the LAN must directly access the WAN communication link. One focus of the NASA Langley Research Center HPCC K-12 Program is the development and evaluation of software for use on local TCP/IP servers that use low-speed connections to minimize direct client WAN access.

The NASA Langley Research Center HPCC K-12 Program has presented a computer infrastructure that is both comprehensive and cost effective. By building a proper foundation, working on a regional basis, and using computer systems based on open standards, a school building can connect to the Internet for $100-150 per month. This figure includes all support and communication links and allows the school building to assign virtually unlimited e-mail accounts. Additionally, the pilot program has demonstrated that dedicated LAN-to-LAN Internet connectivity can be achieved with acceptable performance by using a standard analog phone line when this line is used in conjunction with a distributed server topology.

References

[1]J. Gargano, D. Wasley, *Request for Comments 1709: K-12 Internetworking Guidelines*, (available online at ftp://ds.internic.net/rfc/)

[2]Teach, Carole, ed., *Building the Future: K-12 Network Technology Planning Guide*, California Department of Education, Sacramento, CA, 1994 (available online at gopher://goldmine.cde.ca.us)

Appendix Primary Connection Hub Configuration

A sample configuration of the hardware and software for establishing a connection hub is given.

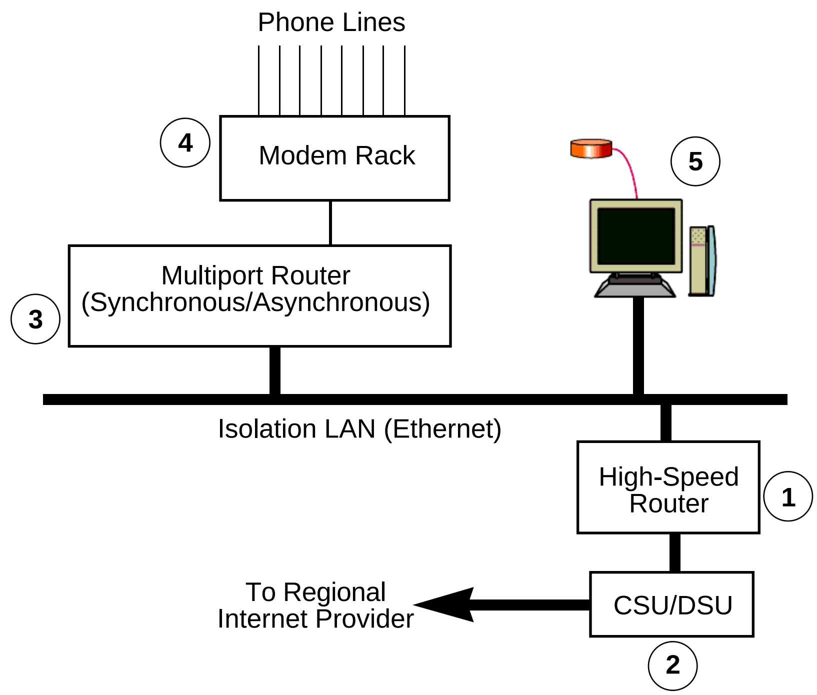

TABLE A1. Hardware for the Isolation LAN

Item	Hardware	Purpose	Example	Estimated Street Price
1.	High-Speed Synchronous Router	Primary router used for connecting to ISP	Cisco	$2500
2.	DSU/CSU	For use with primary router	ADC Kentrox	$1000
3.	Multiport, Low-Speed Asynchronous Router	Connects remote LAN's to isolation LAN	Telebit, Livingston	$8000
4.	Modem Rack Chassis	Houses multiple modem cards	Telebit, US Robotics	$2000
5.	UNIX Workstation	Primary name server for region	Sun, DEC, HP, IBM	$5000

Below is a detailed parts breakdown and an estimated street price for each item used with this architecture. Refer to table 3 in the body of the paper for the specifications of the UNIX workstation.

TABLE A2. Low-Speed Asynchronous Router

Item	Qty	Description	Estimated Street Price
NB 40i Chassis	1	Main CPU and supporting hardware	$3840
4 MB Memory Upgrade	1	Additional memory to support greater number of asynchronous connections	$250
Ethernet Card	1	To allow connection of K-12 ISP's building LAN	$500
MTA Chassis	1	Rackmount chassis and power supply for up to 4 ASYN8(+) modules	$600
ASYN8+	1	Supports 8 async RS-232 ports	$900
ASYN8	3	Supports 8 additional async RS-232 ports (each)	$2700

The Telebit NetBlazer is configured to connect to a bank of modems. The modems used for the HorizonNet testbed are the Telebit 8840R, which cost approximately $850.00.

TABLE A3. Modem Rack Chassis

Item	Qty	Description	Estimated Street Price
T8000 Rack Chassis	2	19 in. rack chassis that supports up to 16 modems, 2 power supplies, and 1 controller	$2200
Power Supplies	2	Provides power to rack	$1600
Telco Cable	2	Connects rack to service provided by local phone company	$100